More on Descriptive Complexity of Second-Order HORN Logics

Shiguang Feng [⋆1] and Xishun Zhao [⋆⋆2]

[1] Institut für Informatik, Universität, Leipzig
D-04109 Leipzig, Germany.
shiguang@informatik.uni-leipzig.de
[2] Institute of Logic and Cognition, Sun Yat-sen University
510275 Guangzhou, P.R.China.
hsszxs@mail.sysu.edu.cn

arXiv:1403.6611v1 [cs.LO] 26 Mar 2014

Abstract. This paper concerns Grädel's question asked in 1992: whether all problems which are in PTIME and closed under substructures are definable in second-order HORN logic (SO-HORN). We introduce revisions of SO-HORN and DATALOG by adding first-order universal quantifiers over the second-order atoms in the bodies of HORN clauses and DATALOG rules. We show that both logics are as expressive as FO(LFP), the least fixed point logic. We also prove that FO(LFP) can not define all of the problems that are in PTIME and closed under substructures. As a corollary, we answer Grädel's question negatively.

Keywords: Second-order HORN, DATALOG, Descriptive Complexity, Closed under substructures, Expressiveness, PTIME

1 Introduction

Descriptive complexity uses logic systems to measure the resources that are necessary to solve a problem in computational complexity. It provides a bridge between complexity theory and mathematical logic. In 1974, Fagin [6] showed in his seminal work that existential second-order logic (∃SO) captures the complexity class NP on finite structures. Up to now, almost every important complexity class has their corresponding capturing logic systems. First-order logic (FO) is a powerful logic and plays a crucial role in classical model theory, but it is very weak in finite structures. By augmenting iterative or recursive operators, many powerful logics are obtained [5]: Deterministic Transitive Closure Logic FO(DTC), Transitive Closure Logic FO(TC), Least Fixed-Point Logic FO(LFP), Inflationary Fixed-Point Logic FO(IFP) and Partial Fixed-Point Logic FO(PFP). FO(LFP) is a fragment of FO(IFP) but they are equally expressive in finite structures. In 1982, Immerman and Vardi [16,11] showed that FO(LFP) captures PTIME on ordered finite structures independently. A central problem in descriptive complexity is whether there is a logic that can capture PTIME effectively on arbitrary finite structures. This question was first raised by Chandra and Harel [3]. It has been shown that if there is no such a logic, then PTIME $\neq$ NP. As a candidate, FO(LFP) lacks the ability to count numbers. It can not define some simple properties, e.g., the parity of a finite set. Immerman [12] first introduced the idea of adding counting quantifier to a logic. Grädel and Otto [10,15] studied FO(IFP,#) logic, which is Inflationary Fixed-Point Logic with counting quantifiers. It is strictly more expressive than FO(IFP)

⋆ The author is supported by *Deutsche Forschungsgemeinschaft* (DFG), GRK 1763 (QuantLA).
⋆⋆ This author is partially supported by the NSFC projects under grant numbers: 61272059, 60970040 and a MOE project under grant number: 11JJD720020.

and can capture PTIME on many classes of graphs. But it is still failed on arbitrary finite structures. A complicated counterexample was found by Cai, Fürer and Immerman [2]. Atserias, Bulatov and Dawar [1] found another problem in PTIME which is more natural, but it is not definable in FO(IFP,#). This problem reveals that FO(IFP,#) lacks the ability to define the rank of a matrix. A. Dawar [4] introduced FO(IFP,rk), which is Inflationary Fixed-point Logic with rank operators. Most of the problems known so far that are not definable in FO(IFP,#) are definable in FO(IFP,rk). But whether FO(IFP,rk) captures PTIME on arbitrary finite structures is still open.

While most of the logics for PTIME are based on fixed-point logic, Grädel [8] introduced second-order Horn logic (SO-HORN), which is a set of second-order formulas whose first-order part is a universal Horn formula with respect to second-order relation variables. Grädel showed that SO-HORN captures PTIME on ordered structures. By the definition, it is easily seen that all SO-HORN formulas are closed under substructures. So if a structure $\mathbf{A}$ satisfies a SO-HORN formula φ, then every substructure of $\mathbf{A}$ also satisfies φ. This property implies that, on arbitrary finite structures, SO-HORN can only define the problems that are closed under substructures, where a problem is closed under substructures if when the answer to an instance of this problem is "Yes", then the answer to every sub-instance of it is also "Yes". For example, the 3-Colorability problem is closed under substructures. An interesting question asked by Grädel is that whether SO-HORN can capture problems that are in PTIME and closed under substructures [9]. A counterpart of SO-HORN is the DATALOG program that originates in database theory. A DATALOG program is a set of rules, where the head of each rule is an atomic formula and the body is a set of atomic or negative atomic formulas. All the relation symbols that occur in the head of some rule are intentional and all other symbols are extensional. The intentional symbols can only occur positively in the body of rule. DATALOG is equivalent to FO(LFP) on ordered finite structures but is strictly less expressive than FO(LFP) on arbitrary finite structures [5].

In this paper, we shall introduce second-order revised Horn logic (SO-HORN^r) and revised DATALOG program (DATALOG^r). In SO-HORN^r formulas, the first-order universal quantifier is allowed to occur over second-order atoms in the body of HORN clauses (treated as implication). Similarly, for a DATALOG^r program we allow universal first-order quantifiers over intentional relations in the body of a rule. We show that both SO-HORN^r and DATALOG^r are strictly more expressive than their original versions, and they are equivalent to FO(LFP) on arbitrary finite structures. Using DATALOG^r we also show that there are problems in PTIME that are closed under substructures but not definable in SO-HORN^r. This implies that FO(LFP) can not define all of the problems that are in PTIME and closed under substructures. As a corollary, we answer the question asked by Grädel negatively.

This paper is organized as follows: In section 2, we at first recall some notations and previous results, then introduce SO-HORN^r and DATALOG^r, and study their expressive powers. We prove that both of them are equivalent to FO(LFP) on arbitrary finite structures. Section 3 is devoted to definitions of encodings of structures. These encodings are important in constructions and proofs in Section 4, in which we show that there are problems in PTIME that are closed under substructures

but not definable in FO(LFP). As a corollary, SO-HORN cannot capture the class of problems that are in PTIME and closed under substructures. We conclude this paper in Section 5.

2 Second-Order HORN Logic, DATALOG and Their Revisions

In this section we shall recall second-order HORN logic and DATALOG program. Then we introduce their revisions and compare their expressiveness with FO(LFP).

2.1 SO-HORN and Its Revision

Definition 1. *[8] Let τ be a vocabulary. Second-order Horn logic over τ, denoted by* SO-HORN(τ)*, is a set of second-order formulas of the form*

$$Q_1 R_1 \cdots Q_m R_m \forall \bar{x}(C_1 \wedge \cdots \wedge C_n)$$

where $Q_i \in \{\forall, \exists\}$, $R_1, \ldots, R_m$ are relation variables, $C_1, \ldots, C_n$ are Horn clauses with respect to $R_1, \ldots, R_m$. More precisely, each C_j is an implication of the form

$$\alpha_1 \wedge \cdots \wedge \alpha_l \wedge \beta_1 \wedge \cdots \wedge \beta_q \rightarrow H$$

where

(1) each α_s is an atomic formula $R_i \bar{x}$,
(2) each β_t is either an atomic formula $P\bar{y}$ or its negation $\neg P\bar{y}$ where $P \in \tau$,
(3) H is either an atomic formula $R_k \bar{u}$ or the Boolean constant $\perp$ (for FALSE).

We use SO-HORN to denote the set all second-order Horn formulas. ($\exists$SO)-HORN is the fragment of SO-HORN where all second-order quantifiers are existential.

Proposition 1. *[8]*

- SO-HORN *captures PTIME on ordered finite structures.*
- SO-HORN *is equivalent to* ($\exists$SO)-HORN*.*

Definition 2. *In Definition 1, if we replace condition (1) by*

(1') each α_s is either an atomic formula $R_i \bar{x}$ or a formula $\forall \bar{z} R_i \bar{y} \bar{z}$,

then we call this logic second-order revised Horn logic, denoted by SO-HORN^r*.*

Example 1. Let $\tau = \{S, f\}$, where S is a unary relation symbol and f is a binary function symbol. Then

$$\exists R \forall x \forall y((Sx \rightarrow Rx)) \wedge (Rx \wedge Ry \rightarrow R(f(x, y))) \wedge (\forall z Rz \rightarrow \perp))$$

is a SO-HORN^r formula.

SO-HORN^r captures PTIME on ordered finite structures [7]. Let ($\exists$SO)-HORN^r denote the existential fragment of SO-HORN^r. We can prove the following proposition using a similar technique in the proof of Theorem 5 of [8].

Proposition 2. SO-HORN^r *is equivalent to* ($\exists$SO)-HORN^r *on arbitrary finite structures.*

Therefore, without loss of generality, we only consider ($\exists$SO)-HORN^r formulas throughout this paper.

2.2 DATALOG and Its Revision

Definition 3. *[5] Let τ be a vocabulary. A* DATALOG *program Π over τ is a finite set of rules of the form*

$$\beta \leftarrow \alpha_1, \ldots, \alpha_l$$

where $l \geq 0$ and

(1) each α_i is either an atomic formula or a negated atomic formula,
(2) β is an atomic formula $R\bar{x}$, where R doesn't occur negated in any rule of Π.

β is the head of the rule and the sequence $\alpha_1, \ldots, \alpha_l$ constitute the body. Every relation symbol occurring in the head of some rule of Π is intentional, and the other symbols in τ are extensional. We use $(\tau, \Pi)_{int}$ and $(\tau, \Pi)_{ext}$ to denote the set of intentional and extensional symbols, respectively.

To compare the expressive power of DATALOG programs with other logics, we allow 0-ary relation symbols. If Q is a 0-ary relation, its value is from $\{\emptyset, \{\emptyset\}\}$. $Q = \emptyset$ means that Q is FALSE and $Q = \{\emptyset\}$ means that Q is TRUE.

In the following we recall the semantics of DATALOG programs. For more details we refer the reader to [5]. Given a DATALOG program Π, to simplify the presentation, we suppose that for any intentional relation symbol P in Π, if P is r-ary and $\bar{x}_P = x_{P,1}, \ldots, x_{P,r}$ are distinct variables, then any rule in Π with head symbol P has the form

$$P\bar{x}_P \leftarrow \alpha_1, \ldots, \alpha_l.$$

Otherwise, we carry out some suitable changes. For example, if $P\mathbf{c} \leftarrow R\mathbf{c}$ is a rule and $\mathbf{c}$ is a constant symbol, we can replace this rule by $Px \leftarrow R\mathbf{c}, x = \mathbf{c}$.

Suppose that $P^1, \ldots, P^m$ are all intentional relation symbols in Π. For each P^i we associate the formula

$$\begin{aligned} \phi_{P^i}(\bar{x}_{P^i}) = \bigvee \{ \exists \bar{v} (\alpha_1 \wedge \cdots \wedge \alpha_l) \mid P^i \bar{x}_{P^i} \leftarrow \alpha_1, \ldots, \alpha_l \in \Pi \\ \text{and } \bar{v} \text{ are the free variables in } \alpha_1 \wedge \cdots \wedge \alpha_l \\ \text{that are different from } \bar{x}_{P^i} \}. \end{aligned} \tag{1}$$

Let $\mathbf{A}$ be a $(\tau, \Pi)_{ext}$-structure and $P^i_{(0)} = \emptyset\, (1 \leq i \leq m)$. If $P^1_{(n)}, \ldots, P^m_{(n)}$ are known and P^i is not 0-ary, we define

$$P^i_{(n+1)} = \{\bar{a} \in A^{\mathrm{arity}(P^i)} \mid (\mathbf{A}, P^1_{(n)}, \ldots, P^m_{(n)}) \models \phi_{P^i}(\bar{x}_{P^i})[\bar{a}]\}.$$

If P^i is 0-ary, we define

$$P^i_{(n+1)} = \{\emptyset \mid (\mathbf{A}, P^1_{(n)}, \ldots, P^m_{(n)}) \models \phi_{P^i}\}.$$

For each $P^i\, (1 \leq i \leq m)$, if $P^i_{(k)} = P^i_{(k+1)}$ for some number k, then define $P^i_{(\infty)} = P^i_{(k)}$. Since $P^1, \ldots, P^m$ occur only positively in each ϕ_{P^i}, the simultaneous fixed-point $(P^1_{(\infty)}, \ldots, P^m_{(\infty)})$ defined by $(\phi_{P^1}, \ldots, \phi_{P^m})$ exists. By applying the rules of Π, every $(\tau, \Pi)_{ext}$-structure $\mathbf{A}$ gives a τ-structure

$$\mathbf{A}[\Pi] = (\mathbf{A}, P^1_{(\infty)}, \ldots, P^m_{(\infty)}).$$

A DATALOG formula has the form $(\Pi, P)\bar{t}$, where P is an r-ary intentional relation symbol and $\bar{t} = t_1, \ldots, t_r$ are variables that don't occur in Π. For a $(\tau, \Pi)_{ext}$-structure $\mathbf{A}$ and $\bar{a} = a_1, \ldots, a_r \in A$,

$$\mathbf{A} \models (\Pi, P)\bar{t}[\bar{a}] \text{ iff } (t_1[\bar{a}], \ldots, t_r[\bar{a}]) \in P_{(\infty)}.$$

If P is 0-ary, then

$$\mathbf{A} \models (\Pi, P) \text{ iff } P_{(\infty)} = \{\emptyset\}.$$

Next we introduce a revision of DATALOG by adding first-order universal quantifiers to it.

Definition 4. *In Definition 3, if each α_i in the body is either an atomic formula or a negated atomic formula, or a formula $\forall\bar{y}R\bar{y}\bar{z}$, where $R \in (\tau, \Pi)_{int}$, then we call this logic program revised* DATALOG *program, denoted by* DATALOG^r *program.*

For a DATALOG^r program Π, the intentional relation symbols also occur only positively in each rule of Π. So the semantics of DATALOG programs can be applied to DATALOG^r programs.

Example 2. The Alternating Graph Accessibility Problem (AGAP):

INPUT: *a directed graph $G=(V, E)$, a partition $V = V_{uni} \cup V_{exi}$ and two nodes $\mathbf{s}, \mathbf{t}$.*
QUESTION: *is there an alternating path from $\mathbf{s}$ to $\mathbf{t}$?*

is a well-known P-complete problem [13]. The nodes in V_{uni} are universal and the nodes in V_{exi} are existential. We say there is an alternating path from x to y if

- *$x = y$, or*
- *$x \in V_{exi}$ and $\exists z \in V$ such that $(x, z) \in E$ and there is an alternating path from z to y, or*
- *$x \in V_{uni}$ and $\forall z \in V$, if $(x, z) \in E$, then there is an alternating path from z to y.*

We can use a DATALOG^r *formula (Π, P) to define this problem, where*

$$\begin{aligned} \Pi : P_{alt}xy &\leftarrow x = y \\ P_{alt}xy &\leftarrow \neg P_{uni}x, Exz, P_{alt}zy \\ P_{alt}xy &\leftarrow P_{uni}x, \forall z Qxzy \\ Qxzy &\leftarrow \neg Exz, y = y \\ Qxzy &\leftarrow P_{alt}zy, x = x \\ P &\leftarrow P_{alt}\mathbf{st} \end{aligned}$$

$(\tau, \Pi)_{int} = \{P_{alt}, Q, P\}$ and $(\tau, \Pi)_{ext} = \{E, P_{uni}, \mathbf{s}, \mathbf{t}\}$. For any $(\tau, \Pi)_{ext}$-structure $\mathbf{A}$, $\mathbf{A} \models (\Pi, P)$ iff there is an alternating path from $\mathbf{s}$ to $\mathbf{t}$.

2.3 Expressiveness of SO-HORN^r and DATALOG^r

In the following we show that each SO-HORN^r sentence is equivalent to the negation of a DATALOG^r formula and vice versa.

Proposition 3. *Let σ be a vocabulary.*

1. *For every sentence ϕ in SO-HORN^r (σ), there exist a* DATALOG^r *program* Π *over the vocabulary* τ *and a formula* (Π, P)*, where* P *is a 0-ary intentional relation symbol, such that* $\sigma = (\tau, \Pi)_{ext}$ *and for any* $(\tau, \Pi)_{ext}$*-structure* $\mathbf{A}$*,*

$$\mathbf{A} \models \phi \text{ iff } \mathbf{A} \not\models (\Pi, P). \tag{2}$$

2. *For every* DATALOG^r *formula* (Π, P)*, where* P *is a 0-ary intentional relation symbol, there exists a* SO-HORN^r *sentence* ϕ *over* $(\tau, \Pi)_{ext}$ *such that the above equivalence (2) also holds.*

Proof. 1. Let ϕ be a SO-HORN^r sentence. Suppose

$$\phi = \exists R_1 \cdots \exists R_m \forall \bar{x}(C_1 \wedge \cdots \wedge C_n),$$

where each C_j is an implication of the form

$$\alpha_1 \wedge \cdots \wedge \alpha_l \rightarrow H.$$

We construct a DATALOG^r program Π as follows:

(1) For each R_i $(1 \leq i \leq m)$, we add $R_i\bar{u} \leftarrow R_i\bar{u}$ as a rule to Π.
(2) For each clause $\alpha_1 \wedge \cdots \wedge \alpha_l \rightarrow H$ in ϕ,
   – if H is an atomic formula $R_i\bar{u}$, then we add the following rule to Π

$$R_i\bar{u} \leftarrow \alpha_1, \ldots, \alpha_l,$$

   – if H is the Boolean constant $\bot$, then we add the following rule to Π

$$P \leftarrow \alpha_1, \ldots, \alpha_l.$$

(3) If no H in the clauses of ϕ is the Boolean constant $\bot$, then we add the following rule to Π

$$P \leftarrow v \neq v,$$

where v is a new variable.

For an arbitrary σ-structure $\mathbf{A}$, to prove $\mathbf{A} \models \phi$ iff $\mathbf{A} \not\models (\Pi, P)$, we first consider the case that no H in the clause of ϕ is the Boolean constant $\bot$. In this case, for each R_i $(1 \leq i \leq m)$, we just let $R_i^A = A^{\text{arity}(R_i)}$, where A is the domain of $\mathbf{A}$. It's easily seen that $(\mathbf{A}, R_1^A, \ldots, R_m^A) \models \forall \bar{x}(C_1 \wedge \cdots \wedge C_n)$, i.e., $\mathbf{A} \models \phi$. Because $P \leftarrow v \neq v$ is the only rule in Π that contains the head symbol P and no structure satisfies $v \neq v$, we have $\mathbf{A} \not\models (\Pi, P)$.

For the other case, we note that $\mathbf{A} \models \phi$ iff there exist minimal relations $R_1^A, \ldots, R_m^A$ such that $(\mathbf{A}, R_1^A, \ldots, R_m^A) \models \forall \bar{x}(C_1 \wedge \cdots \wedge C_n)$ and for any relations $R_1', \ldots, R_m'$, if $(\mathbf{A}, R_1', \ldots, R_m') \models \forall \bar{x}(C_1 \wedge \cdots \wedge C_n)$ then $R_i^A \subseteq R_i'$ $(1 \leq i \leq m)$. If we treat the clauses in ϕ as rules, then the relations $R_1^A, \ldots, R_m^A$ can be computed by the least fixed-point semantics. Remember that we replace every clause $\alpha_1 \wedge \cdots \wedge \alpha_l \rightarrow \bot$ that with head symbol $\bot$ by the rule $P \leftarrow \alpha_1, \ldots, \alpha_l$ in the construction above. It's easily seen that $\mathbf{A} \models \phi$ iff $\mathbf{A} \not\models (\Pi, P)$.

2. To prove the converse direction, we may assume that 0-ary relation symbols only occur in the head of a rule, otherwise, we can carry out some suitable changes.

For example, given a $\mathsf{DATALOG}^r$ formula (Π, Q), we replace Q by $Q'x$ in Π, where Q' and x are new symbols, and add the rule $Q \leftarrow Q'x$ to Π. Let Π' denote the new $\mathsf{DATALOG}^r$ program. It is easily checked that (Π, Q) and (Π', Q) are equivalent.

Let (Π, P) be a $\mathsf{DATALOG}^r$ formula, we construct a $\mathsf{SO\text{-}HORN}^r$ formula ϕ such that (Π, P) and $\neg\phi$ are equivalent. Suppose that $Q_1, \ldots, Q_s$ are all non-0-ary intentional relation symbols and $x_1, \ldots, x_t$ are all free variables in Π, the prefix of ϕ is $\exists Q_1 \cdots \exists Q_s \forall x_1 \cdots \forall x_t$, the matrix of ϕ is obtained by the following constructions.

- For each rule $Q \leftarrow \alpha_1, \ldots, \alpha_l$ in Π, where Q is a 0-ary intentional relation symbol, if Q is the 0-ary intentional relation symbol P, then we add $\alpha_1 \wedge \cdots \wedge \alpha_l \rightarrow \bot$ as a clause, otherwise we ignore the rule.
- For each rule $Q_i\bar{u} \leftarrow \alpha_1, \ldots, \alpha_l$ in Π, where Q_i is a non-0-ary intentional relation symbol, we add $\alpha_1 \wedge \cdots \wedge \alpha_l \rightarrow Q_i\bar{u}$ as a clause.

Similar to the proof above, we can show that for any structure $\mathbf{A}$, $\mathbf{A} \models \phi$ iff $\mathbf{A} \not\models (\Pi, P)$. □

Theorem 1. $\mathsf{DATALOG}^r \equiv \mathsf{FO(LFP)}$ *on arbitrary finite structures.*

Proof. We first show that $\mathsf{DATALOG}^r \preceq$ FO(LFP). Let (Π, P_1) be a $\mathsf{DATALOG}^r$ formula. Suppose that $P_1, \ldots, P_m$ are all intentional relation symbols in Π. Consider the following formula for each P_i.

$$\phi_i(\bar{x}_{P_i}) = \bigvee\{\exists\bar{v}(\alpha_1 \wedge \cdots \wedge \alpha_l) \mid P_i\bar{x}_{P_i} \leftarrow \alpha_1, \ldots, \alpha_l \in \Pi \text{ and } \bar{v} \text{ are the free variables in } \alpha_1 \wedge \cdots \wedge \alpha_l \text{ that are different from } \bar{x}_{P_i}\}$$

Because $P_1, \ldots, P_m$ occur only positively in each ϕ_i $(1 \leq i \leq m)$, $(\phi_1, \ldots, \phi_m)$ defines a least simultaneous fixed-point. The formula (Π, P_1) is equivalent to a FO(S-LFP) formula

$$[\text{S-LFP}_{\bar{x}_1, P_1, \bar{x}_2, P_2, \ldots, \bar{x}_m, P_m}\ \phi_1, \phi_2, \ldots, \phi_m]$$

that is equivalent to a FO(LFP) formula (for details we refer the reader to [5]).

To prove FO(LFP) $\preceq \mathsf{DATALOG}^r$, we use the normal form of FO(LFP). Let u be a variable, write

$$\tilde{u} = \underbrace{u \cdots u}_{i \text{ times}}$$

where i will be clear from the context in the following. Every FO(LFP) formula is equivalent to a formula of the form [5]:

$$\exists u[\text{LFP}_{\bar{z}, Z}\psi(\bar{x}, \bar{z})]\tilde{u}$$

where $\psi(\bar{x}, \bar{z})$ is a first-order formula. Without loss of generality we assume that ψ is in the following DNF form:

$$Q_1 y_1 \cdots Q_m y_m (C_1 \vee \cdots \vee C_n)$$

where $Q_i \in \{\forall, \exists\}$ and each C_j is a conjunction of atomic or negated atomic formulas. Let $\bar{x}, \bar{z}, \bar{y} = y_1, \ldots, y_m$ be the free variables in $C_1 \vee \cdots \vee C_n$. We first construct a $\mathsf{DATALOG}^r$ formula $(\Pi_1, P_1)\bar{t}_1$ that is equivalent to $\psi(\bar{x}, \bar{z})$.

– For each clause $C_j = \alpha_1 \wedge \alpha_2 \wedge \cdots \wedge \alpha_k$, set

$$\Pi_{C_j} = \{P\bar{x}\bar{z}\bar{y} \leftarrow \alpha_1, \alpha_2, \ldots, \alpha_k\}.$$

– Set

$$\Pi_{C_1 \vee \cdots \vee C_n} = \bigcup_{j=1}^{n} \Pi_{C_j}.$$

It is straightforward to prove that $C_1 \vee \cdots \vee C_n(\bar{x}\bar{z}\bar{y})$ and $(\Pi_{C_1 \vee \cdots \vee C_n}, P)\bar{t}$ are equivalent.

– If $Q_m = \forall$, set

$$\Pi_m = \Pi_{C_1 \vee \cdots \vee C_n} \cup \{P_m\bar{x}\bar{z}y_1 \cdots y_{m-1} \leftarrow \forall y_m P\bar{x}\bar{z}\bar{y}\}.$$

– If $Q_m = \exists$, set

$$\Pi_m = \Pi_{C_1 \vee \cdots \vee C_n} \cup \{P_m\bar{x}\bar{z}y_1 \cdots y_{m-1} \leftarrow P\bar{x}\bar{z}\bar{y}\}.$$

Suppose that $Q_k y_k \cdots Q_m y_m(C_1 \vee \cdots \vee C_n)(\bar{x}\bar{z}y_1 \cdots y_{k-1})$ $(1 < k \leq m)$ and $(\Pi_k, P_k)\bar{t}_k$ are equivalent.

– If $Q_{k-1} = \forall$, set

$$\Pi_{k-1} = \Pi_k \cup \{P_{k-1}\bar{x}\bar{z}y_1 \cdots y_{k-2} \leftarrow \forall y_{k-1} P_k\bar{x}\bar{z}y_1 \cdots y_{k-1}\}.$$

– If $Q_{k-1} = \exists$, set

$$\Pi_{k-1} = \Pi_k \cup \{P_{k-1}\bar{x}\bar{z}y_1 \cdots y_{k-2} \leftarrow P_k\bar{x}\bar{z}y_1 \cdots y_{k-1}\}.$$

We repeat these steps above until $(\Pi_1, P_1)\bar{t}_1$ is obtained. It's easily seen that $(\Pi_1, P_1)\bar{t}_1$ and $\psi(\bar{x}, \bar{z})$ are equivalent. Set

$$\Pi = \Pi_1[Z\bar{x}_/Z_] \cup \{Z\bar{x}\bar{z} \leftarrow P_1\bar{x}\bar{z}, Q\bar{x} \leftarrow Z\bar{x}\tilde{u}\}$$

where $\Pi_1[Z\bar{x}_/Z_]$ is obtained by replacing every formula $Z\bar{v}$ by $Z\bar{x}\bar{v}$ in Π_1, and Q is a new relation symbol. By the definition of FO(LFP) we know that Z occurs only positively in each C_j, so Π is a $\mathsf{DATALOG}^r$ program. For any structure $\mathbf{A}$, $\mathbf{A} \models (\Pi, Q)[\bar{a}]$ iff $\mathbf{A} \models \exists u[\mathrm{LFP}_{\bar{z},Z}\psi(\bar{a}/\bar{x}, \bar{z})]\tilde{u}$. □

Because FO(LFP) is closed under negations, combing the theorem above with Proposition 3 we obtain the following theorem.

Theorem 2. $\mathsf{SO\text{-}HORN}^r \equiv \mathsf{FO(LFP)}$ *on arbitrary finite structures.*

3 Encodings of Structures and $\mathsf{DATALOG}^r$

We say that a formula ϕ is closed under substructures, if for any structure $\mathbf{A}$, $\mathbf{A} \models \phi$ implies $\mathbf{B} \models \phi$ for any substructure $\mathbf{B}$ of $\mathbf{A}$. Every SO-HORN formula is closed under substructures. The main task of this paper is to show that there are problems in PTIME that are closed under substructures, but not definable in $\mathsf{SO\text{-}HORN}^r$. The

construction of such a problem involves encodings of structures and $\mathsf{DATALOG}^r$ programs. This section is devoted to the definitions of the encodings which will be used in next section.

In [14], Lindell introduced the notion of invariant relations, which are defined on perfect binary trees and showed that there are problems in PTIME that are not definable in FO(LFP). A perfect binary tree is a full binary tree in which all leaves are at the same level. Let $T = \langle V, E\rangle$ be a perfect binary tree, where V is the set of nodes and E is the set of edges. Suppose R is an r-ary relation on V and f is an automorphism of T. Given a tuple $\bar{a} = (a_1, \ldots, a_r) \in R$, we write $f(\bar{a}) = (f(a_1), \ldots, f(a_r))$ and $f[R] = \{(f(a_1), \ldots, f(a_r)) \mid (a_1, \ldots, a_r) \in R\}$. We say that R is an invariant relation if for every automorphism f, $R = f[R]$. It is easily seen that the equality $=$ and E are invariant relations.

Lemma 1. *If R_1 and R_2 are r-ary invariant relations then $\neg R_1$, $R_1 \cap R_2$ and $R_1 \cup R_2$ are also invariant relations.*

Proof. Let f be an automorphism, it's clear that $\neg f[R_1] = f[\neg R_1]$, $f[R_1] \cap f[R_2] = f[R_1 \cap R_2]$, $f[R_1] \cup f[R_2] = f[R_1 \cup R_2]$. To prove $\neg R_1$ is an invariant relation, we have

$$\begin{aligned}
\bar{a} \in \neg R_1 &\iff \bar{a} \notin R_1 \\
&\overset{(1)}{\iff} \bar{a} \notin f[R_1] \ ((1) : R_1 = f[R_1]) \\
&\iff \bar{a} \in \neg f[R_1] \\
&\overset{(2)}{\iff} \bar{a} \in f[\neg R_1] \ ((2) : \neg f[R_1] = f[\neg R_1]).
\end{aligned}$$

To prove $R_1 \cap R_2$ is an invariant relation, we have

$$\begin{aligned}
\bar{a} \in R_1 \cap R_2 &\iff \bar{a} \in R_1 \text{ and } \bar{a} \in R_2 \\
&\overset{(1)}{\iff} \bar{a} \in f[R_1] \text{ and } \bar{a} \in f[R_2] \\
&\qquad ((1) : R_1 = f[R_1], R_2 = f[R_2]) \\
&\iff \bar{a} \in f[R_1] \cap f[R_2] \\
&\overset{(2)}{\iff} \bar{a} \in f[R_1 \cap R_2] \\
&\qquad ((2) : f[R_1] \cap f[R_2] = f[R_1 \cap R_2]).
\end{aligned}$$

To prove $R_1 \cup R_2$ is an invariant relation, we have

$$\begin{aligned}
\bar{a} \in R_1 \cup R_2 &\iff \bar{a} \in R_1 \text{ or } \bar{a} \in R_2 \\
&\overset{(1)}{\iff} \bar{a} \in f[R_1] \text{ or } \bar{a} \in f[R_2] \\
&\qquad ((1) : R_1 = f[R_1], R_2 = f[R_2]) \\
&\iff \bar{a} \in f[R_1] \cup f[R_2] \\
&\overset{(2)}{\iff} \bar{a} \in f[R_1 \cup R_2] \\
&\qquad ((2) : f[R_1] \cup f[R_2] = f[R_1 \cup R_2]).
\end{aligned}$$

□

Lemma 2. *Suppose R is an r-ary invariant relation, R' is a k-ary invariant relation and g is a permutation of $\{1, \ldots, r\}$. Define*

$$\begin{aligned}
R_1 &= \{(a_{g(1)}, \ldots, a_{g(r)}) \mid (a_1, \ldots, a_r) \in R\} \\
R_2 &= \{(a_1, \ldots, a_r, b_1, \ldots, b_k) \mid (a_1, \ldots, a_r) \in R \text{ and } (b_1, \ldots, b_k) \in R'\}.
\end{aligned}$$

Then R_1 and R_2 are also invariant relations.

Lemma 3. *If R is a $(r+k)$-ary invariant relation. Define*

$$R_1 = \{(a_1, \ldots, a_r) \mid (b_1, \ldots, b_k, a_1, \ldots, a_r) \in R \text{ for all nodes } b_1, \ldots, b_k\}$$
$$R_2 = \{(a_1, \ldots, a_r) \mid \exists b_1, \ldots, b_k \text{ such that } (b_1, \ldots, b_k, a_1, \ldots, a_r) \in R\}.$$

Then R_1 and R_2 are also invariant relations.

Proof. For any automorphism f we show that $R_1 = f[R_1]$ and $R_2 = f[R_2]$. To prove $R_1 = f[R_1]$, we have

$$\begin{aligned}
\bar{a} \in R_1 &\overset{(1)}{\Longleftrightarrow} \bar{b}\bar{a} \in R \text{ for any tuple } \bar{b} \ ((1) : \text{by the definition of } R_1) \\
&\overset{(2)}{\Longleftrightarrow} \bar{b}\bar{a} \in f[R] \text{ for any tuple } \bar{b} \ ((2) : R = f[R]) \\
&\overset{(3)}{\Longleftrightarrow} \bar{b}\bar{a}' \in R \text{ for any tuple } \bar{b}, \text{ where} \\
&\qquad f(\bar{a}') = \bar{a} \ ((3) : f \text{ is an automorphism}) \\
&\overset{(1)}{\Longleftrightarrow} \bar{a}' \in R_1 \\
&\overset{(4)}{\Longleftrightarrow} \bar{a} \in f[R_1] \ ((4) : f(\bar{a}') = \bar{a}).
\end{aligned}$$

To prove $R_2 = f[R_2]$, we have

$$\begin{aligned}
\bar{a} \in R_2 &\overset{(1)}{\Longleftrightarrow} \exists \bar{b} \text{ such that } \bar{b}\bar{a} \in R \ ((1) : \text{by the definition of } R_2) \\
&\overset{(2)}{\Longleftrightarrow} \exists \bar{b} \text{ such that } \bar{b}\bar{a} \in f[R] \ ((2) : R = f[R]) \\
&\overset{(3)}{\Longleftrightarrow} \exists \bar{b} \text{ such that } \bar{b}'\bar{a}' \in R, \text{ where } f(\bar{b}') = \bar{b}, \\
&\qquad f(\bar{a}') = \bar{a} \ ((3) : f \text{ is an automorphism}) \\
&\overset{(1)}{\Longleftrightarrow} \bar{a}' \in R_2 \\
&\overset{(4)}{\Longleftrightarrow} \bar{a} \in f[R_2] \ ((4) : f(\bar{a}') = \bar{a}).
\end{aligned}$$

□

Let a, b be two nodes of a perfect binary tree T, we use $a \barwedge b$ and $d(a)$ to denote the least common ancestor of a, b and the depth of a, respectively. For example, in the following perfect binary tree

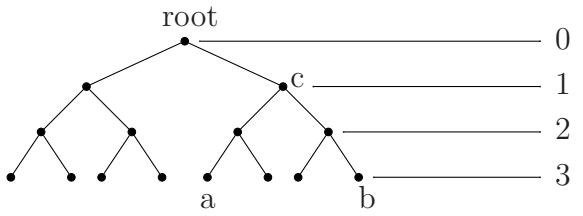

we have $d(\mathbf{root}) = 0$, $d(c) = 1$, $d(a) = d(b) = 3$, and $a \barwedge b = c$.

Let $(a_1, \ldots, a_r)$ be an r-ary tuple of nodes, its characteristic tuple is defined as

$$\begin{aligned}
(a_1, \ldots, a_r)^* = (&d(a_1), d(a_1 \barwedge a_2), \ldots, d(a_1 \barwedge a_r), \\
&d(a_2), d(a_2 \barwedge a_3), \ldots, d(a_2 \barwedge a_r), \\
&\ldots, d(a_r))
\end{aligned}$$

which is a $\frac{r(r+1)}{2}$-ary tuple of numbers.

Let R be an invariant relation of T, the characteristic relation of R is defined to be

$$R^* = \{(a_1, \ldots, a_r)^* \mid (a_1, \ldots, a_r) \in R\}.$$

The following propositions were shown in [14]:

- For any tuples $(a_1, \ldots, a_r)$ and $(b_1, \ldots, b_r)$, we have that $(a_1, \ldots, a_r)^* = (b_1, \ldots, b_r)^*$ iff there is an automorphism f of T such that $f(a_i) = b_i\,(1 \leq i \leq r)$,
- For any two invariant relations R_1 and R_2, we have $R_1 = R_2$ iff $R_1^* = R_2^*$.
- If R is an r-ary invariant relation and $(a_1, \ldots, a_r)^* = (b_1, \ldots, b_r)^*$, then $(a_1, \ldots, a_r) \in R$ iff $(b_1, \ldots, b_r) \in R$.

Next we introduce the notion of saturated relations and use it for the definition of encodings.

Definition 5. *Let T be a perfect binary tree, R an r-ary relation on T. We say that R is a saturated relation on T if for any nodes $a_1, \ldots, a_r$, $b_1, \ldots, b_r$ of T, whenever $d(a_i) = d(b_i)\,(1 \leq i \leq r)$ we have*

$$(a_1, \ldots, a_r) \in R \text{ iff } (b_1, \ldots, b_r) \in R.$$

It is clear that a saturated relation is also invariant.

In the following we fix a vocabulary $\tau = \{R_1, \ldots, R_k\}$, and write $\tau' = \tau \cup \{\mathbf{root}, E\}$, where **root** is a constant symbol and E is a binary relation symbol. We define a class of τ'-structures

$$\mathcal{T} = \{\langle V, \mathbf{root}, E, R_1, \ldots, R_k\rangle \mid \langle V, E\rangle \text{ is a perfect binary tree, } \mathbf{root} \text{ is the root node and } R_1, \ldots, R_k \text{ are saturated relations on } \langle V, E\rangle\}.$$

Definition 6. *Let $\mathbf{A} = \langle\{0, 1, \ldots, h-1\}, R_1^A, \ldots, R_k^A\rangle$ be an arbitrary τ-structure, and let $T = \langle V, E\rangle$ be a perfect binary tree with depth h. We can expand T into a τ'-structure $C(\mathbf{A}, T) = \langle V, \mathbf{root}, E, R_1^T, \ldots, R_k^T\rangle \in \mathcal{T}$ as follows: For any $R_i\,(1 \leq i \leq k)$ and any nodes $a_1, \ldots, a_{r_i} \in V$,*

$$C(\mathbf{A}, T) \models R_i^T a_1 \cdots a_{r_i} \text{ iff } \mathbf{A} \models R_i^A d(a_1) \cdots d(a_{r_i}),$$

where r_i is the arity of R_i, and $d(a_j)\,(1 \leq j \leq r_i)$ is the depth of a_j.

Roughly speaking, $C(\mathbf{A}, T)$ contains the same information as $\mathbf{A}$, but the size of $C(\mathbf{A}, T)$ is exponential in the size of $\mathbf{A}$.

Definition 7. *Suppose $\mathbf{T} = \langle V, \mathbf{root}, E, R_1^T, \ldots, R_k^T\rangle \in \mathcal{T}$. We define a τ-structure $\mathbf{A}$ as follows: (1) The domain is $\{0, \ldots, h-1\}$ with h the depth of $\mathbf{T}$, and (2) For each $i = 1, \ldots, k$,*

$$R_i^A = \{(n_1, \ldots, n_{r_i}) \mid \exists a_1, \ldots, a_{r_i} \in \mathbf{T} \text{ such that } R_i^T a_1 \cdots a_{r_i} \text{ and } n_1 = d(a_1), \ldots, n_{r_i} = d(a_{r_i})\}$$

where r_i is the arity of R_i. It is not hard to see that $\mathbf{T} = C(\mathbf{A}, (V, E))$. So we denote $\mathbf{A}$ by $C^{-1}(\mathbf{T})$.

For a fixed natural number m we introduce a relation symbol FUL_m with arity m. Let

$$\sigma = \{R_1^*, \ldots, R_k^*, (\neg R_1)^*, \ldots, (\neg R_k)^*\},$$
$$\sigma' = \sigma \cup \{\mathbf{0}, SUCC, R_{\neq}, R_{\neg e}, FUL_m^*\}.$$

Definition 8. *For each $\mathbf{T} = \langle V, \mathbf{root}, E, R_1, \ldots, R_k\rangle \in \mathcal{T}$, we define a σ'-structure*

$$S_{\mathbf{T}} = \Big\langle \{0, 1, \ldots, h-1\}, \mathbf{0}, SUCC, R_{\neq}, R_{\neg e}, FUL_m^*, R_1^*, \ldots, R_k^*, (\neg R_1)^*, \ldots, (\neg R_k)^* \Big\rangle,$$

where h is the depth of $\mathbf{T}$, $\mathbf{0}$ is a constant interpreted by 0, $SUCC$ is the successor relation on the domain, FUL_m is the full relation $\{0, \ldots, h-1\}^m$, and $R_{\neq}, R_{\neg e}, FUL_m^*$, $R_1^*, \ldots, R_k^*$, $(\neg R_1)^*, \ldots, (\neg R_k)^*$ are the characteristic relations of $\neq, (\neg E)$, $FUL_m, R_1, \ldots, R_k$, $\neg R_1, \ldots, \neg R_k$, respectively.

Let $\Pi = \{\gamma_1, \ldots, \gamma_s\}$ be a $\mathsf{DATALOG}^r$ program on $\mathcal{T}$. Suppose $X_1, \ldots, X_w$ are all intentional relation symbols in Π and for each rule γ_i, let n_{γ_i} be the number of free variables that occur in γ_i. Define

$$m = \max\{n_{\gamma_1}, \ldots, n_{\gamma_s}, \mathrm{arity}(R_1), \ldots, \mathrm{arity}(R_k), \mathrm{arity}(X_1), \ldots, \mathrm{arity}(X_w)\}.$$

Next we shall construct, based on Π, a $\mathsf{DATALOG}^r$ program Π^* such that for any $\mathsf{DATALOG}^r$ formula (Π, P), there is a corresponding $\mathsf{DATALOG}^r$ formula (Π^*, P^*) such that $\mathbf{T} \models (\Pi, P)$ iff $S_{\mathbf{T}} \models (\Pi^*, P^*)$ for any $\mathbf{T} \in \mathcal{T}$, where P and P^* are 0-ary intentional relation symbols.

Every element of $\mathbf{T}$ is a node of a perfect binary tree, while every element of $S_{\mathbf{T}}$ is a number which can be treat as the depth of some node. So for each variable x in Π, we introduce a new variable i_x, and for any two variables x_1, x_2 in Π, we introduce a new variable $i_{x_1 \bar{\wedge} x_2}$.

For $\bar{x} = x_1 \cdots x_r$, we shall use the following abbreviations:

$$(\bar{x})^* = i_{x_1} i_{x_1 \bar{\wedge} x_2} \cdots i_{x_1 \bar{\wedge} x_r} i_{x_2} \cdots i_{x_{r-1} \bar{\wedge} x_r} i_{x_r}$$
$$\forall(\bar{x})^* = \forall i_{x_1} \forall i_{x_1 \bar{\wedge} x_2} \cdots \forall i_{x_1 \bar{\wedge} x_r} \forall i_{x_2} \cdots \forall i_{x_{r-1} \bar{\wedge} x_r} \forall i_{x_r}.$$

From now on, for some technical reason, we treat $i_{u \bar{\wedge} v}$ and $i_{v \bar{\wedge} u}$ as the same variable.

Now we begin our construction of Π^*.

For each rule $\beta \leftarrow \alpha_1, \ldots, \alpha_l$ in Π, suppose that $v_1, \ldots, v_n$ are the free variables in it, we add formula

$$FUL_m v_1 v_2 \cdots v_{n-1} v_n v_n \cdots v_n$$

to the body and obtain a new rule

$$\beta \leftarrow \alpha_1, \ldots, \alpha_l, FUL_m v_1 v_2 \cdots v_{n-1} v_n \cdots v_n.$$

For each new rule, we

- replace $x = y$ by $i_x = i_{x \bar{\wedge} y}, i_{x \bar{\wedge} y} = i_y$ (reason: $d(x) = d(x \mathbin{\bar{\wedge}} y) = d(y)$), for constant **root**, we replace $i_{\mathbf{root}}$ by constant $\mathbf{0}$, and replace $i_{\mathbf{root} \bar{\wedge} x}$ also by $\mathbf{0}$, since $\mathbf{root} \mathbin{\bar{\wedge}} a = \mathbf{root}$ for any node a.

- replace Exy by $i_x = i_{x\bar{\wedge}y}, SUCCi_{x\bar{\wedge}y}i_y$ (reason: $d(x) = d(x\ \bar{\wedge}\ y)$, $d(x\ \bar{\wedge}\ y) + 1 = d(y)$).
- replace $x \neq y$ by $R_{\neq}i_x i_{x\bar{\wedge}y} i_y$ (reason: $R_{\neq}$ is the characteristic relation of $\neq$).
- replace $\neg Exy$ by $R_{\neg e}i_x i_{x\bar{\wedge}y} i_y$ (reason: $R_{\neg e}$ is the characteristic relation of $\neg E$).
- replace $R\bar{x}$ by $R^*(\bar{x})^*$, where R is a symbol in $\{R_1, \ldots, R_k, FUL_m\}$, or an intentional relation symbol which is not 0-ary.
- replace $\neg R\bar{x}$ by $(\neg R)^*(\bar{x})^*$, where R is a symbol in $\{R_1, \ldots, R_k\}$.
- replace $\forall y_1 \cdots \forall y_t R y_1 \cdots y_t z_1 \cdots z_s$ by

$$\Phi_R = \begin{pmatrix} FUL_m^*(z_1 z_2 \cdots z_{s-1} z_s \cdots z_s)^* \wedge \\ \forall (y_1 \cdots y_t)^* \forall i_{y_1 \bar{\wedge} z_1} \cdots \forall i_{y_1 \bar{\wedge} z_s} \forall i_{y_2 \bar{\wedge} z_1} \cdots \forall i_{y_t \bar{\wedge} z_{s-1}} \forall i_{y_t \bar{\wedge} z_s} \\ (FUL_m^*(y_1 \cdots y_t z_1 \cdots z_s z_s \cdots z_s)^* \rightarrow R^*(y_1 \cdots y_t z_1 \cdots z_s)^*) \end{pmatrix}$$

  where R is an intentional relation symbol such that $\text{arity}(R) > 0$.
- replace P by P^*, where P is a 0-ary intentional relation symbol.

We denote the resulting program by Π^*. Please note that Π^* is not a $\mathsf{DATALOG}^r$ program since Φ_R is neither an atom nor a negated atom. Fortunately, Φ_R is equivalent to a $\mathsf{DATALOG}^r$ formula $(\Pi_1, Q_2)\bar{t}$, where

$$\begin{aligned} \Pi_1 : Q(y_1 \cdots y_t z_1 \cdots z_s)^* &\leftarrow \neg FUL_m^*(y_1 \cdots y_t z_1 \cdots z_s z_s \cdots z_s)^* \\ Q(y_1 \cdots y_t z_1 \cdots z_s)^* &\leftarrow R^*(y_1 \cdots y_t z_1 \cdots z_s)^* \\ Q_1(z_1 \cdots z_s)^* &\leftarrow \forall (y_1 \cdots y_t)^* \forall i_{y_1 \bar{\wedge} z_1} \cdots \forall i_{y_1 \bar{\wedge} z_s} \\ &\qquad \forall i_{y_2 \bar{\wedge} z_1} \cdots \forall i_{y_t \bar{\wedge} z_{s-1}} \forall i_{y_t \bar{\wedge} z_s} Q(y_1 \cdots y_t z_1 \cdots z_s)^* \\ Q_2(z_1 \cdots z_s)^* &\leftarrow Q_1(z_1 \cdots z_s)^*, FUL_m^*(z_1 z_2 \cdots z_{s-1} z_s \cdots z_s)^*. \end{aligned}$$

Now we add Π_1 to Π^* and replace Φ_R by $Q_2(z_1 \cdots z_s)^*$. Then we get a $\mathsf{DATALOG}^r$ program Π_2. It is not hard to see that for every common intentional relation symbol X of Π^* and Π_2, $(\Pi^*, X)\bar{t}'$ and $(\Pi_2, X)\bar{t}'$ are equivalent. Therefore, without loss of generality, we shall use Π^* directly in future proofs.

Remark 1. The reason that we replace the formula $\phi_R = \forall y_1 \cdots \forall y_t R y_1 \cdots y_t z_1 \cdots z_s$ by Φ_R rather than

$$\begin{aligned} \forall (y_1 \cdots y_t)^* \forall i_{y_1 \bar{\wedge} z_1} \cdots \forall i_{y_1 \bar{\wedge} z_s} \forall i_{y_2 \bar{\wedge} z_1} \cdots \forall i_{y_t \bar{\wedge} z_{s-1}} \forall i_{y_t \bar{\wedge} z_s} \\ R^*(y_1 \cdots y_t z_1 \cdots z_s)^* \end{aligned}$$

is that there exist $\mathbf{T} \in \mathcal{T}$, $\bar{a} \in \mathbf{T}$, and an invariant relation R such that

$$\mathbf{T} \models \forall y_1 \cdots \forall y_t R y_1 \cdots y_t z_1 \cdots z_s[\bar{a}],$$

but

$$\begin{aligned} S_{\mathbf{T}} \not\models \forall (y_1 \cdots y_t)^* \forall i_{y_1 \bar{\wedge} z_1} \cdots \forall i_{y_1 \bar{\wedge} z_s} \forall i_{y_2 \bar{\wedge} z_1} \cdots \forall i_{y_t \bar{\wedge} z_{s-1}} \forall i_{y_t \bar{\wedge} z_s} \\ R^*(y_1 \cdots y_t z_1 \cdots z_s)^*[(\bar{a})^*]. \end{aligned}$$

It is easy to check that ϕ_R is equivalent to the formula

$$\begin{aligned} &FUL_m z_1 \cdots z_{s-1} z_s z_s \cdots z_s \wedge \\ &\forall y_1 \cdots \forall y_t (FUL_m y_1 \cdots y_t z_1 \cdots z_s z_s \cdots z_s \rightarrow R y_1 \cdots y_t z_1 \cdots z_s). \end{aligned}$$

We replace FUL_m and R by their characteristic relations FUL_m^* and R^*, respectively, then we can obtain Φ_R. The formula

$$FUL_m^*(z_1 z_2 \cdots z_{s-1} z_s \cdots z_s)^*$$

guarantees that if there exists some $\bar{e} \in S_{\mathbf{T}}$ such that $S_{\mathbf{T}} \models \Phi_R[\bar{e}]$, then $\bar{e}$ is the characteristic tuple of some tuple $\bar{a}$ of $\mathbf{T}$, i.e. $\bar{e} = (\bar{a})^*$.

Lemma 4. *Given a formula $\phi_R = \forall y_1 \cdots \forall y_t R y_1 \cdots y_t z_1 \cdots z_s$, a structure $\mathbf{T} \in \mathcal{T}$. Let Φ_R be defined as above. Define $Q_1 = \{\bar{a} \mid \mathbf{T} \models \phi_R[\bar{a}]\}$, $Q_2 = \{\bar{e} \mid S_{\mathbf{T}} \models \Phi_R[\bar{e}]\}$. Suppose R is an invariant relation on $\mathbf{T}$. Then $(Q_1)^* = Q_2$.*

Proof. Because R is an invariant relation, by the definition of Q_1 and Lemma 1 and 3, we know that Q_1 is also an invariant relation. We first show that $(Q_1)^* \subseteq Q_2$. Suppose $\bar{e} \in (Q_1)^*$ for some $\bar{e} \in S_{\mathbf{T}}$, there must exist a tuple $\bar{a}$ of $\mathbf{T}$ such that $\bar{a} \in Q_1$, $(\bar{a})^* = \bar{e}$, and $\bar{b}\bar{a} \in R$ for all tuples $\bar{b}$ of $\mathbf{T}$, i.e.,

$$\begin{aligned}\mathbf{T} \models \Big(&FUL_m z_1 \cdots z_{s-1} z_s z_s \cdots z_s \wedge \\ &\forall y_1 \cdots \forall y_t (FUL_m y_1 \cdots y_t z_1 \cdots z_s z_s \cdots z_s \rightarrow R y_1 \cdots y_t z_1 \cdots z_s)\Big)[\bar{a}].\end{aligned}$$

By the definition of Φ_R we can see that $S_{\mathbf{T}} \models \Phi_R[(\bar{a})^*]$, which implies $\bar{e} \in Q_2$.

Next we show that $Q_2 \subseteq (Q_1)^*$. Consider an arbitrary $\bar{e} \in S_{\mathbf{T}}$ such that $\bar{e} \in Q_2$. By the definition of Q_2 and Φ_R, we have $S_{\mathbf{T}} \models FUL_m(z_1 \cdots z_s z_s \cdots z_s)^*[\bar{e}]$, so there exists a tuple $\bar{a}$ of $\mathbf{T}$ such that $\bar{e} = (\bar{a})^*$. On the contrary, suppose that $\bar{a} \notin Q_1$, then there exists some tuple $\bar{b}$ such that $\bar{b}\bar{a} \notin R$. Because R is an invariant relation, for any tuples $\bar{b}'$ and $\bar{a}'$, if $(\bar{b}\bar{a})^* = (\bar{b}'\bar{a}')^*$ then $\bar{b}'\bar{a}' \notin R$. Combing that R^* is the characteristic relation of R we know that

$$S_{\mathbf{T}} \not\models R^*(y_1 \cdots y_t z_1 \cdots z_s)^*[(\bar{b}\bar{a})^*] \tag{3}$$

and

$$S_{\mathbf{T}} \models FUL_m^*(y_1 \cdots y_t z_1 \cdots z_s z_s \cdots z_s)^*[(\bar{b}\bar{a})^*]. \tag{4}$$

(3) and (4) give $S_{\mathbf{T}} \not\models \Phi_R[(\bar{a})^*]$, i.e., $\bar{e} \notin Q_2$, contrary to the assumption that $\bar{e} \in Q_2$. So $\bar{a} \in Q_1$, which implies $\bar{e} \in (Q_1)^*$. □

We assume that a 0-ary relation is also an invariant relation. The characteristic relation of a 0-ary relation is $\emptyset^* = \emptyset$ and $\{\emptyset\}^* = \{\emptyset\}$.

Proposition 4. *For any intentional relation symbol P in Π and any $\mathbf{T} \in \mathcal{T}$, $P^{\mathbf{T}[\Pi]}$ is an invariant relation on $\mathbf{T}$ and $(P^{\mathbf{T}[\Pi]})^* = (P^*)^{S_{\mathbf{T}}[\Pi^*]}$. Moreover, if P is a 0-ary intentional relation symbol, then $\mathbf{T} \models (\Pi, P)$ iff $S_{\mathbf{T}} \models (\Pi^*, P^*)$.*

Proof. We first show that if P is an intentional relation symbol in Π and $\mathbf{T}$ is a structure in $\mathcal{T}$, then $P^{\mathbf{T}[\Pi]}$ is an invariant relation on $\mathbf{T}$. Let $P^1, \ldots, P^{m'}$ be all intentional relation symbols in Π. Consider the following formula constructed for each P^i (see (1))

$$\begin{aligned}\phi_{P^i}(\bar{x}_{P^i}) = \bigvee\{\exists \bar{v}(\alpha_1 \wedge \cdots \wedge \alpha_l) \mid\ & P^i \bar{x}_{P^i} \leftarrow \alpha_1, \ldots, \alpha_l \in \Pi \\ & \text{and } \bar{v} \text{ are the free variables in } \alpha_1 \wedge \cdots \wedge \alpha_l \\ & \text{that are different from } \bar{x}_{P^i}\}.\end{aligned}$$

If the relation defined by each α_s is an invariant relation, then by Lemma 1, 2 and 3 we can see that the relation defined by ϕ_{P^i} is also an invariant relation. Each α_s is either an (negated) atomic formula with relation symbol in

$$\{=, E, R_1, \ldots, R_k, \neq, \neg E, \neg R_1, \ldots, \neg R_k\}$$

where the relations defined by them are all invariant relations, or an atomic formula $P^j\bar{x}$ or a formula $\forall\bar{y}P^j\bar{y}\bar{z}\,(1 \leq j \leq m')$. When computing the fixed-point of $P^1, \ldots, P^{m'}$, we set $P^i_{(0)} = \emptyset\,(1 \leq i \leq m')$, where $\emptyset$ is an invariant relation. In Lemma 3 we showed that if P^j is an invariant relation then the relation defined by $\forall\bar{y}P^j\bar{y}\bar{z}$ is also an invariant relation. If $P^1_{(n)}, \ldots, P^{m'}_{(n)}$ are invariant relations, then each

$$P^i_{(n+1)} = \{\bar{a} \mid (\mathbf{T}, P^1_{(n)}, \ldots, P^{m'}_{(n)}) \models \phi_{P^i}(\bar{x}_{P^i})[\bar{a}]\}$$

or

$$P^i_{(n+1)} = \{\emptyset \mid (\mathbf{T}, P^1_{(n)}, \ldots, P^{m'}_{(n)}) \models \phi_{P^i}\}$$

is also an invariant relation. So the fixed-points $P^1_{(\infty)}, \ldots, P^{m'}_{(\infty)}$ are invariant relations, i.e., $P^{\mathbf{T}[\Pi]}$ is an invariant relation on $\mathbf{T}$.

Next we show that $(P^{\mathbf{T}[\Pi]})^* = (P^*)^{S_{\mathbf{T}}[\Pi^*]}$. We first prove $(P^i_{(n)})^* = (P^i)^*_{(n)}\,(1 \leq i \leq m')$ for each $n \geq 0$ by induction.

Basis: If $n = 0$, then $P^i_{(0)} = \emptyset$, $(P^i)^*_{(0)} = \emptyset\,(1 \leq i \leq m')$. We have $(P^i_{(0)})^* = (P^i)^*_{(0)}\,(1 \leq i \leq m')$.

Inductive step: Assuming $(P^i_{(k)})^* = (P^i)^*_{(k)}\,(1 \leq i \leq m')$, we show that $(P^i_{(k+1)})^* = (P^i)^*_{(k+1)}\,(1 \leq i \leq m')$. The case where P^i is 0-ary is easy, in the following we only consider the case where P^i is not 0-ary.

To prove $(P^i_{(k+1)})^* \subseteq (P^i)^*_{(k+1)}$, suppose $\bar{e} \in (P^i_{(k+1)})^*$ for some $\bar{e} \in S_{\mathbf{T}}$, there must exist a tuple $\bar{a}$ of $\mathbf{T}$ such that $\bar{a} \in P^i_{(k+1)}$ and $\bar{e} = (\bar{a})^*$. By the semantics of $\mathsf{DATALOG}^r$ we know that

$$P^i_{(k+1)} = \{\bar{a} \mid (\mathbf{T}, P^1_{(k)}, \ldots, P^{m'}_{(k)}) \models \phi_{P^i}(\bar{x}_{P^i})[\bar{a}]\}.$$

By the definition of ϕ_{P^i}, there exist a rule $P^i\bar{x}_{P^i} \leftarrow \alpha_1, \ldots, \alpha_l$ in Π such that

$$(\mathbf{T}, P^1_{(k)}, \ldots, P^{m'}_{(k)}) \models \exists\bar{v}(\alpha_1 \wedge \cdots \wedge \alpha_l)[\bar{a}],$$

therefore there exists some $\bar{b}$ such that

$$(\mathbf{T}, P^1_{(k)}, \ldots, P^{m'}_{(k)}) \models (\alpha_1 \wedge \cdots \wedge \alpha_l)[\bar{a}\bar{b}].$$

Because $P^i\bar{x}_{P^i} \leftarrow \alpha_1, \ldots, \alpha_l$ is a rule of Π, we can infer that

$$(P^i)^*(\bar{x}_{P^i})^* \leftarrow \alpha'_1, \ldots, \alpha'_l, FUL^*_m(\bar{x}_{P^i}\bar{v}\tilde{v}')^*$$

is a rule of Π^*, where $\alpha'_1, \ldots, \alpha'_l$ and $FUL^*_m(\bar{x}_{P^i}\bar{v}\tilde{v}')^*$ are obtained by replacing $\alpha_1, \ldots, \alpha_l, FUL_m$ by the corresponding formulas in the construction of Π^*. Remember that we replace $\forall\bar{y}R\bar{y}\bar{z}$ by Φ_R and in Lemma 4 we showed that the relation defined by Φ_R is the characteristic relation of the relation defined by $\forall\bar{y}R\bar{y}\bar{z}$. By

the definition of $S_{\mathbf{T}}$ and the induction hypothesis $(P^i_{(k)})^* = (P^i)^*_{(k)}\,(1 \le i \le m')$ we know that

$$(S_{\mathbf{T}}, (P^1)^*_{(k)}, \ldots, (P^{m'})^*_{(k)}) \models (\alpha'_1 \wedge \cdots \wedge \alpha'_l \wedge FUL^*_m(\bar{x}_{P^i}\bar{v}\tilde{v}')^*)[(\bar{a}\bar{b})^*],$$

i.e.,

$$(S_{\mathbf{T}}, (P^1)^*_{(k)}, \ldots, (P^{m'})^*_{(k)}) \models \exists\bar{u}(\alpha'_1 \wedge \cdots \wedge \alpha'_l \wedge FUL^*_m(\bar{x}_{P^i}\bar{v}\tilde{v}')^*)[(\bar{a})^*]$$

where $\bar{u}$ are the free variables in $\alpha'_1 \wedge \cdots \wedge \alpha'_l \wedge FUL^*_m(\bar{x}_{P^i}\bar{v}\tilde{v}')^*$ that are different from $(\bar{x}_{P^i})^*$. Combining $\bar{e} = (\bar{a})^*$ we obtain $\bar{e} \in (P^i)^*_{(k+1)}$.

Next we show that $(P^i)^*_{(k+1)} \subseteq (P^i_{(k+1)})^*$. Suppose $\bar{e} \in (P^i)^*_{(k+1)}$ for some $\bar{e} \in S_{\mathbf{T}}$. There must exist a rule

$$(P^i)^*(\bar{x}_{P^i})^* \leftarrow \alpha'_1, \ldots, \alpha'_l, FUL^*_m(\bar{x}_{P^i}\bar{v}\tilde{v}')^* \tag{5}$$

in Π^* such that

$$(S_{\mathbf{T}}, (P^1)^*_{(k)}, \ldots, (P^{m'})^*_{(k)}) \models \exists\bar{u}(\alpha'_1 \wedge \cdots \wedge \alpha'_l \wedge FUL^*_m(\bar{x}_{P^i}\bar{v}\tilde{v}')^*)[\bar{e}].$$

Hence there exists a tuple $\bar{f} \in S_A$ such that

$$(S_{\mathbf{T}}, (P^1)^*_{(k)}, \ldots, (P^{m'})^*_{(k)}) \models (\alpha'_1 \wedge \cdots \wedge \alpha'_l \wedge FUL^*_m(\bar{x}_{P^i}\bar{v}\tilde{v}')^*)[\bar{e}\bar{f}].$$

The formula $FUL^*_m(\bar{x}_{P^i}\bar{v}\tilde{v}')^*$ guarantees that $(\bar{a}\bar{b})^* = \bar{e}\bar{f}$ and $(\bar{a})^* = \bar{e}$ for some tuple $\bar{a}\bar{b}$ of $\mathbf{T}$. Because $P^1_{(k)}, \ldots, P^{m'}_{(k)}$ are invariant relations, by the induction hypothesis $(P^i_{(k)})^* = (P^i)^*_{(k)}\,(1 \le i \le m')$ we know that

$$(\mathbf{T}, P^1_{(k)}, \ldots, P^{m'}_{(k)}) \models (\alpha_1 \wedge \cdots \wedge \alpha_l)[\bar{a}\bar{b}]$$

where $P^i\bar{x}_{P^i} \leftarrow \alpha_1, \ldots, \alpha_l$ is the corresponding rule of (5) in Π. So $\bar{a} \in P^i_{(k+1)}$. Combining $\bar{e} = (\bar{a})^*$ we obtain $\bar{e} \in (P^i_{(k+1)})^*$.

By the proof above we know that $(P^{\mathbf{T}[\Pi]})^* = (P^*)^{S_{\mathbf{T}}[\Pi^*]}$. If P is a 0-ary intentional relation symbol, then $\mathbf{T} \models (\Pi, P)$ iff $S_{\mathbf{T}} \models (\Pi^*, P^*)$. □

4 On Grädel's Question

The complexity class EXPTIME is the set of all decision problems that are decidable by a deterministic Turing machine in $O(2^{n^c})$ time for some constant number c, where n is the length of the input. By Time Hierarchy Theorem, we know that PTIME $\subset$ EXPTIME, i.e., PTIME is a proper subset of EXPTIME.

Same as the last section, we fix a vocabulary $\tau = \{R_1, \ldots, R_k\}$. Consider an arbitrary class $\mathcal{K}$ of τ-structures decidable deterministically in time 2^{n^c}. Next we shall construct a class of τ'-structures which is in PTIME and closed under substructures. The class $\{\mathbf{T} \in \mathcal{T} \mid C^{-1}(\mathbf{T}) \in \mathcal{K}\}$ might be a candidate (here $\mathcal{T}$ is defined just before Definition 6, and C^{-1} is defined in Definition 7). However, since the substructure of a perfect binary tree is not necessarily perfect, the candidate class is not closed under substructure. Next we introduce some techniques to overcome it.

Definition 9. *Let c be a fixed natural number and $\mathbf{A} = \langle \{0, \ldots, h-1\}, R_1^A, \ldots, R_k^A \rangle$ be a τ-structure. The trivial extension $\mathbf{A}$ is the structure*

$$\langle \{0, \ldots, h + h^c - 1\}, R_1^A, \ldots, R_k^A \rangle,$$

denoted by $\mathbf{A}^+$.

Informally speaking, $\mathbf{A}^+$ is obtained from $\mathbf{A}$ by extending just the domain and keeping all relations unchanged.

For a technical reason we introduce a new unary relation symbol P and let $\tau' = \tau \cup \{\mathbf{root}, E, P\}$. From now on when we speak of a τ'-structure $\mathbf{G} = \langle V, \mathbf{root}, E, P, R_1, \cdots, R_k \rangle$, we assume that

(1) (V, E) is a directed acyclic graph and the nodes reachable from **root** form a binary tree, and
(2) all relations $P, R_1, \cdots, R_k$ when restricted to $T(\mathbf{G})$ are saturated, where $T(\mathbf{G})$ is the largest perfect binary subtree of $\mathbf{G}$ with **root** as the root.

Please note that if a τ'-structure $\mathbf{G}$ satisfies the condition (1) and (2) above, then all its substructures also satisfy these two conditions.

Definition 10. *We define based on $\mathcal{K}$ a class $\mathcal{K}'$ of τ'-structures as follows. Given any $\mathbf{G} = \langle V, \mathbf{root}, E, P, R_1, \ldots, R_k \rangle$, let h be the largest number such that all nodes in the first h levels are marked by P. $\mathbf{G} \in \mathcal{K}'$ iff the following Condition 1 or Condition 2 holds:*

Condition 1:
(1) The depth of $T(\mathbf{G})$ is $h + h^c$.
(2) No node in the last h^c levels is marked by P.
(3) $C^{-1}(T(\mathbf{G}))$ is the trivial extension of $C^{-1}(T_h(\mathbf{G}))$, where $T_h(\mathbf{G})$ is the substructure obtained from $T(\mathbf{G})$ by restricting to the first h levels. That is, for any R_i and any tuple $(a_1, \ldots, a_{r_i})$ containing a node in the last h^c levels, $R_i a_1 \cdots a_{r_i}$ does not hold, where r_i is the arity of R_i.
(4) $C^{-1}(T_h(\mathbf{G})) \in \mathcal{K}$.
Condition 2: *The depth of $T(\mathbf{G})$ is less than $h + h^c$.*

Proposition 5.

1. *$\mathcal{K}'$ is closed under substructures*
2. *$\mathcal{K}' \in$ PTIME, that is, the problem of determining whether a given τ'-structure belongs to $\mathcal{K}'$ can be decided deterministically in polynomial time.*

Proof. 1. Suppose $\mathbf{G}$ is a τ'-structure in $\mathcal{K}'$, then it satisfies either Condition 1 or Condition 2 of Definition 10. Consider an arbitrary substructure $\mathbf{H}$ of $\mathbf{G}$. If $\mathbf{G}$ satisfies Condition 2, it's clear that $\mathbf{H}$ also satisfies Condition 2, and hence it is in $\mathcal{K}'$. Suppose $\mathbf{G}$ satisfies Condition 1, then the perfect binary tree $T(\mathbf{H})$ either equals to $T(\mathbf{G})$, where in this case $\mathbf{H}$ satisfies Condition 1, or the depth of $T(\mathbf{H})$ is less than that of $T(\mathbf{G})$, where in this case $\mathbf{H}$ satisfies Condition 2. Altogether, $\mathbf{H} \in \mathcal{K}'$.

2. Consider an arbitrary τ'-structure $\mathbf{G}$. To check that $\mathbf{G} \in \mathcal{K}'$ we need to do the following job.

(1) Check that $\langle V, E\rangle$ is a directed acyclic graph.
(2) Check that all nodes reachable from **root** form a binary tree.
(3) Compute $T(\mathbf{G})$, the largest perfect binary subtree with **root** as root.
(4) Check that the relations $P, R_1, \ldots, R_k$ are saturated on $T(\mathbf{G})$.
(5) Compute the largest number h such that all nodes in the first h levels have property P and check that no nodes in the last h^c levels has property P.
(6) If the depth of $T(\mathbf{G})$ is $h + h^c$ then
  (6.1) Check that (1), (2) and (3) in Condition 1 hold.
  (6.2) Check that $C^{-1}(T_h(\mathbf{G})) \in \mathcal{K}$.
(7) Check that the depth of $T(\mathbf{G})$ is less than $h + h^c$.

It is not hard to see that all tasks except (6.2) can be done in polynomial time. Please note that $\mathcal{K}$ is decidable in 2^{h^c} time and $C^{-1}(T_h(\mathbf{G}))$ has h elements. Task (6.2) can be done in time 2^{h^c}. Please note that if the depth of $T(\mathbf{G})$ is $h+h^c$ then the input size is at least 2^{h+h^c}. Therefore, task (6.2) can be done in polynomial time. □

Our aim is to prove that there is class of τ'-structures in PTIME that is closed under substructures, but not definable in $\mathsf{DATALOG}^r$. Suppose otherwise, then there is a $\mathsf{DATALOG}^r$ formula (Π, Q) such that the class of all finite models of (Π, Q) coincides with $\mathcal{K}'$, where Q is a 0-ary intentional relation symbol.

Consider an arbitrary τ-structure $\mathbf{A} = \langle\{0, \ldots, h-1\}, R_1^A, \ldots, R_1^A\rangle$. Technically we mark every element of $\mathbf{A}$ by P, that is, we interpret P as the full domain of $\mathbf{A}$. For convenience we still denote the new structure by $\mathbf{A}$. Let $\mathbf{A}^+$ be the trivial extension of $\mathbf{A}$, that is, $\mathbf{A}^+$ is obtained from $\mathbf{A}$ just by extending the domain to $\{0, \ldots, h + h^c - 1\}$. Let $\mathbf{T}$ be a τ'-structure in $\mathcal{T}$ such that $C^{-1}(\mathbf{T}) = \mathbf{A}$. Then obviously $\mathbf{T}$ satisfies (1), (2) and (3) in Condition 1. Thus, we have

$$\mathbf{A} \in \mathcal{K} \quad \text{iff} \quad \mathbf{T} \in \mathcal{K}' \ \text{ iff } \ \mathbf{T} \models (\Pi, Q) \tag{6}$$

Let

$$\sigma = \{\mathbf{0}, SUCC, R_{\neq}, R_{\neg e}, FUL_m^*, P^*, R_1^*, \ldots, R_k^*, (\neg P)^*, (\neg R_1)^*, \ldots, (\neg R_k)^*\}$$

be a vocabulary. By Definition 8 we can compute a σ-structure $S_{\mathbf{T}}$ from $\mathbf{T}$:

$$S_{\mathbf{T}} = \big\langle\{0, 1, \ldots, h + h^c - 1\}, \mathbf{0}, SUCC, R_{\neq}, R_{\neg e}, FUL_m^*, P^*, \\ R_1^*, \ldots, R_k^*, (\neg P)^*, (\neg R_1)^*, \ldots, (\neg R_k)^*\big\rangle$$

where $\mathbf{0}$ is a constant interpreted by 0, $SUCC$ is the successor relation on $\{0, 1, \ldots, h+h^c-1\}$ and $R_{\neq}, R_{\neg e}$, FUL_m^*, P^*, $R_1^*, \ldots, R_k^*, (\neg P)^*, (\neg R_1)^*, \ldots, (\neg R_k)^*$ are the characteristic relations of $\neq, (\neg E)$, FUL_m, P, $R_1, \ldots, R_k$, $\neg P, \neg R_1, \ldots, \neg R_k$, respectively.

Let $\mathbf{T}, \mathbf{T}' \in \mathcal{T}$ be such that $C^{-1}(\mathbf{T}) = C^{-1}(\mathbf{T}') = \mathbf{A}^+$. Then $\mathbf{T}$ and $\mathbf{T}'$ must be isomorphic, and hence $S_{\mathbf{T}} = S_{\mathbf{T}'}$. That is, $S_{\mathbf{T}}$ is uniquely determined by $\mathbf{A}$. So, we use $S_{\mathbf{A}}$ to denote $S_{\mathbf{T}}$. The next proposition shows that $S_{\mathbf{A}}$ can be computed efficiently directly from $\mathbf{A}$.

Proposition 6. *$S_{\mathbf{A}}$ is logspace computable from $\mathbf{A}$.*

Before we prove Proposition 6, we show how a contradiction can be derived. By the encodings and Proposition 4 in Section 3, we know that there is another DATALOG^r formula (Π^*, Q^*) such that $\mathbf{T} \models (\Pi, Q)$ iff $S_{\mathbf{T}} \models (\Pi^*, Q^*)$ for any $\mathbf{T} \in \mathcal{T}$. Combining equivalence (6) we have that $\mathbf{A} \in \mathcal{K}$ iff $S_{\mathbf{A}} \models (\Pi^*, Q^*)$ for any τ-structure $\mathbf{A}$.

By Theorem 1, DATALOG^r has the same expressive power as FO(LFP). Thus it is decidable in polynomial time whether a structure is a model of (Π^*, Q^*). Then it follows from Proposition 6 that $\mathcal{K} \in$ PTIME. This would imply EXPTIME=PTIME, since $\mathcal{K}$ is an arbitrary problem in EXPTIME. Because PTIME≠EXPTIME, we have

Proposition 7. *There is a problem in* PTIME *that is closed under substructures but neither definable in* DATALOG^r *nor in* SO-HORN^r

Because SO-HORN $\subset$ SO-HORN$^r \equiv$ FO(LFP), we get the following results.

Theorem 3. FO(LFP) *cannot define all of the problems that are in* PTIME *and closed under substructures.*

Corollary 1. SO-HORN *cannot capture the class of problems that are in* PTIME *and closed under substructures.*

Our Final task is to prove Proposition 6.

Proof of Proposition 6:
Given a τ-structure

$$\mathbf{A} = \langle \{0, 1, 2, \ldots, h-1\}, R_1, \ldots, R_k \rangle ,$$

the corresponding σ-structure is:

$$S_{\mathbf{A}} = \big\langle \{0, 1, \ldots, h + h^c - 1\}, \mathbf{0}, SUCC, R_{\neq}, R_{\neg e}, FUL^*_m, P^*, \\ R^*_1, \ldots, R^*_k, (\neg P)^*, (\neg R_1)^*, \ldots, (\neg R_k)^* \big\rangle.$$

Let $\mathbf{A}$ be the input. It is easily seen that $\{0, 1, \ldots, h + h^c - 1\}, \mathbf{0}, SUCC$ are computable in logspace. By the definition of $S_{\mathbf{A}}$, we know that

$$P^* = \{0, 1, \ldots, h-1\}$$

$$(\neg P)^* = \{h, h+1, \ldots, h + h^c - 1\},$$

so P^* and $(\neg P)^*$ are also computable in logspace.

For the other relations, it is clear that they are computable in logspace if we can decide them in logspace.

Consider the relation $R_{\neq}$, for any tuple (e_1, e_2, e_3), $(e_1, e_2, e_3) \in R_{\neq}$ iff (e_1, e_2, e_3) is the characteristic tuple $(d(a), d(a \barwedge b), d(b))$ of some tuple (a, b) of nodes of $\mathbf{T}$ such that $C^{-1}(\mathbf{T}) = \mathbf{A}^+$ and $a \neq b$. If $d(a) = d(b)$, then a, b are in the same level, hence $d(a \barwedge b) < d(a)$. If $d(a) \neq d(b)$, then a, b are not in the same level, and $d(a \barwedge b) \leq \min\{d(a), d(b)\}$. Conversely, if ($e_1 = e_3$ and $e_2 < e_1$) or ($e_1 \neq e_3$ and $e_2 \leq \min\{e_1, e_3\}$), then there must be some nodes a, b of $\mathbf{T}$ such that $a \neq b$ and $(d(a), d(a \barwedge b), d(b)) = (e_1, e_2, e_3)$. Thus, we can decide $R_{\neq}$ by the following procedure.

INPUT: a tuple (e_1, e_2, e_3)
OUTPUT: accept if $(e_1, e_2, e_3) \in R_{\neq}$, and reject otherwise
1 if $e_1 = e_3$ **then**
  if $e_2 < e_1$ **then** accept **else** reject
2 if $e_1 \neq e_3$ **then**
  if $e_2 \leq \min\{e_1, e_3\}$ **then** accept **else** reject
END

For the relation $R_{\neg e}$, $(e_1, e_2, e_3) \in R_{\neg e}$ iff (e_1, e_2, e_3) is the characteristic tuple $(d(a), d(a \barwedge b), d(b))$ of some tuple (a, b) of nodes of $\mathbf{T}$ such that $C^{-1}(\mathbf{T}) = \mathbf{A}^+$ and $\neg Eab$. Please note that Eab iff $d(a) + 1 = d(b)$ and $d(a \barwedge b) = d(a)$. Thus, if $d(a) + 1 = d(b)$ then $\neg Eab$ iff $d(a \barwedge b) < d(a)$, and if $d(b) \leq d(a)$ or $d(a) < d(b) - 1$ then $\neg Eab$. Please note that $(e_1, e_2, e_3) \in R_{\neg e}$ implies that (e_1, e_2, e_3) is the character tuple of some (a, b) of $\mathbf{T}$. Therefore, we can decide $R_{\neg e}$ by the following procedure.

INPUT: a tuple (e_1, e_2, e_3)
OUTPUT: accept if $(e_1, e_2, e_3) \in R_{\neg e}$, and reject otherwise
1 if $e_1 \geq e_3$ **then**
  if $e_2 \leq e_3$ **then** accept **else** reject
2 if $e_1 < e_3$ **then**
  if $e_1 + 1 = e_3$ **then**
    if $e_2 < e_1$ **then** accept **else** reject
  if $e_1 + 1 < e_3$ **then**
    if $e_2 \leq e_1$ **then** accept **else** reject
END

Suppose Q is a relation in $\{FUL_m^*, R_1^*, \ldots, R_k^*, (\neg R_1)^*, \ldots, (\neg R_k)^*\}$. If a tuple $(e_1, e_{1\barwedge 2}, \ldots, e_{n-1\barwedge n}, e_n) \in Q$, then it must be the characteristic tuple $(d(a_1), d(a_1 \barwedge a_2), \ldots, d(a_{n-1} \barwedge a_n), d(a_n))$ of some tuple $(a_1, a_2, \ldots, a_n)$ of nodes of $\mathbf{T}$ such that $C^{-1}(\mathbf{T}) = \mathbf{A}^+$. Next we show how to decide whether a tuple is a characteristic tuple.

Let $A = \{a_1, \ldots, a_n\}$ be a set of nodes of $\mathbf{T}$ and b be the least common ancestor of the nodes in A. Define

$$\begin{aligned}
\kappa_1 &= \min\{d(a_1), \ldots, d(a_n)\} \\
\kappa_2 &= \min\{d(a_i \barwedge a_j) \mid 1 \leq i < j \leq n\} \\
A_l &= \{a \in A \mid a \text{ is in the left subtree of } b\} \\
A_r &= \{a \in A \mid a \text{ is in the right subtree of } b\}
\end{aligned}$$

where the left (resp. right) subtree is the subtree with the left (resp. right) child node of b as root. It is easily seen that the following conditions must be satisfied.

1. $\kappa_1 \geq \kappa_2$ and $\kappa_2 = d(b)$.
2. If $\kappa_1 = \kappa_2$, then $b \in A$, $\{b\} \cup A_l \cup A_r = A$ and
   - for any $a_i \in A$, $d(b) = d(a_i)$ implies $b = a_i$
   - for any $a_i \in A_l \cup A_r$, $d(b \barwedge a_i) = \kappa_1$
   - for any $a_i, a_j \in A_l$, $d(a_i \barwedge a_j) > \kappa_1$
   - for any $a_i, a_j \in A_r$, $d(a_i \barwedge a_j) > \kappa_1$

- for any $a_i \in A_l$ and $a_j \in A_r$, $d(a_i \,\overline{\wedge}\, a_j) = \kappa_1$

3. If $\kappa_1 > \kappa_2$, then $b \notin A$, $A_l \cup A_r = A$ and
   - for any $a_i, a_j \in A_l$, $d(a_i \,\overline{\wedge}\, a_j) > \kappa_2$
   - for any $a_i, a_j \in A_r$, $d(a_i \,\overline{\wedge}\, a_j) > \kappa_2$
   - for any $a_i \in A_l$ and $a_j \in A_r$, $d(a_i \,\overline{\wedge}\, a_j) = \kappa_2$

Next we present a procedure PreCHECK to check these conditions using the tuple $(d(a_1), d(a_1 \,\overline{\wedge}\, a_2), \ldots, d(a_{n-1} \,\overline{\wedge}\, a_n), d(a_n))$ instead of the nodes $a_1, \ldots, a_n$. Let $(e_1, e_{1\bar{\wedge}2}, \ldots, e_{n-1\bar{\wedge}n}, e_n)$ be an input. If it is a characteristic tuple of some tuple of nodes, then PreCHECK must accept it. But if PreCHECK accepts, it is not necessarily a characteristic tuple.

PROCEDURE PreCHECK
INPUT: a tuple $(e_1, e_{1\bar{\wedge}2}, \ldots, e_{n-1\bar{\wedge}n}, e_n)$
OUTPUT: accept if this tuple satisfies the conditions above, and reject otherwise

1 let κ_1 be the minimal value in set $\{e_1, e_2 \ldots, e_n\}$ and κ_2 be the minimal value in set $\{e_{i\bar{\wedge}j} \mid 1 \le i < j \le n\}$.
2 **if** $\kappa_1 < \kappa_2$ **then** reject
3 **if** $\kappa_1 = \kappa_2$, let e_s be the first element that is not equal to κ_1 in the tuple $(e_1, e_2 \ldots, e_n)$, and define the following three sets

$$S = \{e_i \mid 1 \le i \le n \text{ and } e_i = \kappa_1\}$$
$$S_l = \{e_i \mid e_i \in \{e_1, e_2 \ldots, e_n\}/S \text{ and } e_{s\bar{\wedge}i} > \kappa_1\} \cup \{e_s\}$$
$$S_r = \{e_i \mid e_i \in \{e_1, e_2 \ldots, e_n\}/S \text{ and } e_{s\bar{\wedge}i} = \kappa_1\}$$

then reject if one of the following conditions is satisfied
- $\exists\, e_i, e_j \in S$ such that $e_{i\bar{\wedge}j} \neq \kappa_1$ (or $e_{j\bar{\wedge}i} \neq \kappa_1$)
- $\exists\, e_i \in S$ and $e_j \in S_l \cup S_r$ such that $e_{i\bar{\wedge}j} \neq \kappa_1$ (or $e_{j\bar{\wedge}i} \neq \kappa_1$)
- $\exists\, e_i, e_j \in S_l$ such that $e_{i\bar{\wedge}j} = \kappa_1$ (or $e_{j\bar{\wedge}i} = \kappa_1$)
- $\exists\, e_i, e_j \in S_r$ such that $e_{i\bar{\wedge}j} = \kappa_1$ (or $e_{j\bar{\wedge}i} = \kappa_1$)
- $\exists\, e_i \in S_l$ and $e_j \in S_r$ such that $e_{i\bar{\wedge}j} \neq \kappa_1$ (or $e_{j\bar{\wedge}i} \neq \kappa_1$)

else accept
4 **if** $\kappa_1 > \kappa_2$, define the following two sets

$$S_l = \{e_i \mid 1 < i \le n \text{ and } e_{1\bar{\wedge}i} > \kappa_2\} \cup \{e_1\}$$
$$S_r = \{e_i \mid 1 < i \le n \text{ and } e_{1\bar{\wedge}i} = \kappa_2\}$$

then reject if one of the following conditions is satisfied
- $\exists\, e_i, e_j \in S_l$ such that $e_{i\bar{\wedge}j} = \kappa_2$ (or $e_{j\bar{\wedge}i} = \kappa_2$)
- $\exists\, e_i, e_j \in S_r$ such that $e_{i\bar{\wedge}j} = \kappa_2$ (or $e_{j\bar{\wedge}i} = \kappa_2$)
- $\exists\, e_i \in S_l$ and $e_j \in S_r$ such that $e_{i\bar{\wedge}j} \neq \kappa_2$ (or $e_{j\bar{\wedge}i} \neq \kappa_2$)

else accept
END

The following algorithm can check whether a tuple $(e_1, e_{1\bar{\wedge}2}, \ldots, e_{n-1\bar{\wedge}n}, e_n)$ is a characteristic tuple by recursion.

PROCEDURE CHECK
INPUT: a tuple $(e_1, e_{1\bar{\wedge}2}, \ldots, e_{n-1\bar{\wedge}n}, e_n)$
OUTPUT: accept if $\bar{e}$ is a characteristic tuple of some tuple of nodes, and reject otherwise
1 if $|\bar{e}| = 0$ or $|\bar{e}| = 1$ **then** accept
2 if $|\bar{e}| > 1$ **then** rejects if PreCHECK($\bar{e}$) reject **else** let κ_1 be the minimal value in set $\{e_1, \ldots, e_n\}$ and κ_2 be the minimal value in set $\{e_{i\bar{\wedge}j} \mid 1 \leq i < j \leq n\}$
2.1 if $\kappa_1 = \kappa_2$, let e_s be the first element that is not equal to κ_1 in the tuple $(e_1, \ldots, e_n)$, and define the following three sets

$$S = \{e_i \mid 1 \leq i \leq n \text{ and } e_i = \kappa_1\}$$
$$S_l = \{e_i \mid e_i \in \{e_1, \ldots, e_n\}/S \text{ and } e_{s\bar{\wedge}i} > \kappa_1\} \cup \{e_s\}$$
$$S_r = \{e_i \mid e_i \in \{e_1, \ldots, e_n\}/S \text{ and } e_{s\bar{\wedge}i} = \kappa_1\}$$

and define two tuples

$$\bar{e}_l = (e_{l_1}, e_{l_1\bar{\wedge}l_2}, \ldots, e_{l_{h-1}}, e_{l_{h-1}\bar{\wedge}l_h}, e_{l_h})$$
$$\bar{e}_r = (e_{r_1}, e_{r_1\bar{\wedge}r_2}, \ldots, e_{r_{g-1}}, e_{r_{g-1}\bar{\wedge}r_g}, e_{r_g})$$

where $\{e_{l_1}, \ldots, e_{l_{h-1}}, e_{l_h}\} = S_l$ and $\{e_{r_1}, \ldots, e_{r_{g-1}}, e_{r_g}\} = S_r$
then reject if CHECK($\bar{e}_l$) reject **else**
if CHECK($\bar{e}_r$) reject **then** reject **else** accept
2.2 if $\kappa_1 > \kappa_2$, define two sets

$$S_l = \{e_i \mid 1 < i \leq n \text{ and } e_{1\bar{\wedge}i} > \kappa_2\} \cup \{e_1\}$$
$$S_r = \{e_i \mid 1 < i \leq n \text{ and } e_{1\bar{\wedge}i} = \kappa_2\}$$

and define two tuples

$$\bar{e}_l = (e_{l_1}, e_{l_1\bar{\wedge}l_2}, \ldots, e_{l_{h-1}}, e_{l_{h-1}\bar{\wedge}l_h}, e_{l_h})$$
$$\bar{e}_r = (e_{r_1}, e_{r_1\bar{\wedge}r_2}, \ldots, e_{r_{g-1}}, e_{r_{g-1}\bar{\wedge}r_g}, e_{r_g})$$

where $\{e_{l_1}, \ldots, e_{l_{h-1}}, e_{l_h}\} = S_l$ and $\{e_{r_1}, \ldots, e_{r_{g-1}}, e_{r_g}\} = S_r$
then reject if CHECK($\bar{e}_l$) reject **else**
if CHECK($\bar{e}_r$) reject **then** reject **else** accept
END

Finally, we show how to decide the relation R_1^* as an example, the others are the same as it. Suppose the corresponding relation R_1 is r-ary.

INPUT: a tuple $(e_1, e_{1\bar{\wedge}2}, \ldots, e_{r-1\bar{\wedge}r}, e_r)$
OUTPUT: accept if $\bar{e} \in R_1^*$, and reject otherwise
if CHECK($\bar{e}$) reject **then** reject **else**
if $(e_1, e_2, \ldots, e_r) \in R_1$ **then** accept **else** reject
END

In the above algorithm, the recursion depth of the procedure CHECK is at most r, and the procedure PreCHECK only uses logspace with respect to the length of the input structure, so this algorithm finishes in logspace.

Altogether we complete the proof.

5 Conclusion

In this paper, we have introduced the SO-HORN^r logic and the DATALOG^r program, and show that both of them are equivalent to FO(LFP) on arbitrary structures. We also show that FO(LFP) can not define all of the problems that are in PTIME and closed under substructures. As a corollary, SO-HORN cannot capture the class of problems that are in PTIME and closed under substructures. But it is still open that whether every problem that are in SO-HORN^r and closed under substructures are definable in SO-HORN.

References

1. A. Atserias, A. Bulatov, and A. Dawar. Affine systems of equations and counting infinitary logic. Theoretical Computer Science, 410:1666-1683,2009.
2. J. Cai, M. Fürer, N. Immerman. An Optimal Lower Bound on the Number of Variables for Graph Identification. Combinatorica, 1992, 12(4): 389-410.
3. A. Chandra, D. Harel. Structure and Complexity of Relational Queries. J. Comp. Syst. Sciences, 1982, 25: 99-128.
4. A. Dawar, M. Grohe, B. Holm, and B. Laubner. Logics with rank operators. In Proceedings of the 24th IEEE Symposium on Logic in Computer Science, pages 113-122,2009.
5. H.-D. Ebbinghaus, J. Flum. Finite Model Theory. Springer, 1999.
6. R. Fagin. Generalized First-order Spectra and Polynomial-Time Recognizable Sets. In R.M. Karp(ed.), Complexity of Computation, SIAM-AMS Proceedings, 1974, 7: 43-73.
7. S. Feng and X. Zhao. Complexity and Expressive Power of Second-Order Extended Horn Logic, Mathematical Logic Quarterly, 59(1-2): 4-11, 2013.
8. E. Grdäel. The Expressive Power of Second-order Horn Logic. Proceedings of the 8th annual symposium on Theoretical aspects of computer science, p.466-477, February 1991.
9. E. Grädel. Capturing Complexity Classes by Fragments of Second-order Logic. Theoretical Computer Science, 101: 35-57, 1992.
10. E. Grädel, M. Otto. Inductive Definability with Counting on Finite Structures. In E. Borger et al. (eds.), CSL '92, Lecture Notes in Computer Science, 1993, 702: 231-247.
11. N. Immerman. Relational Queries Computable in Polynomial Time. Information and Control, 1986, 68: 86-104. Extended abstract in Proc. 14th Annual ACM Symp. on Theory of Computing, 1982: 147-152.
12. N. Immerman. Expressibility as a Complexity Measure: results and directions. In Second Structure in Complexity Conference, IEEE Computer Society Press, 1987, pages: 194-202.
13. N. Immerman. Number of quantifiers is better than number of tape cells. Journal of Computer and System Sciences, vol. 22, no. 3, 1981: 384-406.
14. Steven Lindell. An Analysis of Fixed-point Queries on Binary Trees. Theoretical Computer Science, Volume 85, Issue 1, 5 August 1991, Pages 75-95.
15. M. Otto. The Expressive Power of Fixed-point Logic with Counting. Journal of Symbolic Logic, 1996, 61: 147-176.
16. M. Vardi. The Complexity of Relational Query Languages. Proc. of 14th STOC, 1982: 137-146.